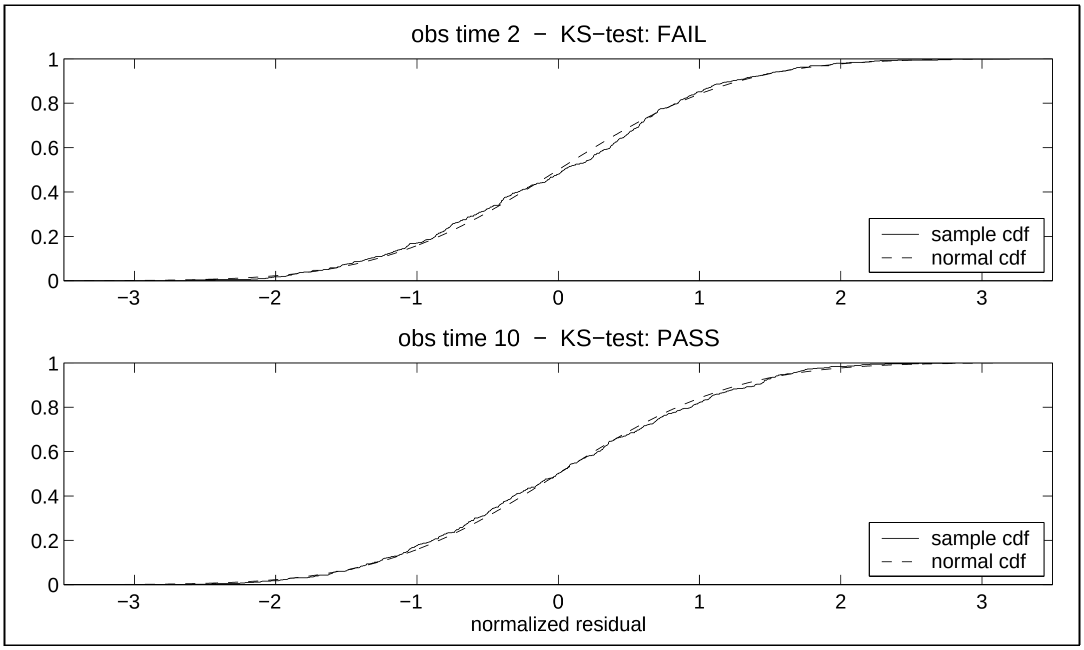

Figure 6.9: Sample cumulative distribution function (cdf) of the standardized posterior data residuals for Reference Experiment I for two of the twelve brightness images that have been assimilated (solid line). Also shown is the theoretical cumulative distribution function of the standard normal distribution (dashed line). For both images, the sample cdf is close to normal. However, the residuals at the second observation time do not pass the Kolmogorov-Smirnov test for normality at a 5% significance level. The data residuals for one other observation time out of twelve also fail this test. We attribute the deviation from normality to the nonlinearities in the hydrologic model and the measurement process.

6.1.6 Adjoint Variables

Figure 6.10 illustrates the temporal behavior of the adjoint variables. Time series of the adjoint saturation and the adjoint soil temperature are shown for a representative pixel. For comparison, we also plot the time series of the estimated brightness temperature together with the (synthetic) observations. Note that in this experiment the observation pixels and the estimation pixels coincide.

The forcing of the adjoint equation only consists of impulses at the observation times (2.15). These impulses are proportional to the misfits between the observed and the estimated brightness temperature. They are also proportional to the sensitivity derivative of the brightness temperature with respect to the corresponding state variable. Starting from a final condition of zero at the end of the assimilation interval (day 183), the adjoint equation is integrated backward in time. The adjoint therefore equals zero until, going backwards in time, we hit the last measurement time (day 181.6). At the last measurement time, the misfit between the observed and the estimated brightness is negative for this particular pixel. Since the soil saturation is inversely correlated with the brightness temperature, that is $\partial T_B / \partial W_g \leq 0$, the negative data misfit leads to a positive impulse for the adjoint saturation. The opposite is true for the adjoint soil temperature. Further integration backward in time causes the adjoint states to decay and grow exponentially according to the model physics while they are repeatedly forced with data misfit impulses at observation times.

In Figure 6.10 we can also observe how the brightness misfit forcing affects the saturation profile. Since the brightness temperature directly depends on the saturation in the top two nodes at $0cm$ and $-5cm$ via the top layer microwave emissivity and via the heat capacity, the misfit forcing impulse leads to instantaneous effects in the corresponding adjoint saturation components. In contrast, the lower components of the saturation do not directly influence the brightness temperature. Therefore, the adjoint saturation of the lower nodes only experience a delayed effect after the misfit forcing has been propagated downward through the (adjoint) model physics.

Finally, Figure 6.10 also illustrates the much longer memory of soil moisture compared to soil temperature. After each measurement time, the adjoint soil temperature decays (backwards) to zero within twelve hours, whereas the adjoint saturation does not reach zero before the misfit impulse at the next earlier measurement time comes in.

6.2 Reference Experiment II

To complement Reference Experiment I of Section 6.1, we now present another synthetic experiment using different error statistics for the uncertain parameters and a different realization, that is a different seed for the random number generator. The focus of Reference Experiment II is on the model error. Unless otherwise stated, the inputs for this experiment are the same as for Reference Experiment I. The initial condition is still uncertain, but with somewhat reduced variance compared to Reference Experiment I. For this experiment, the initial top node saturation varies between 0.3 and 0.65 (Figure 6.11). The sample mean is 0.52 and the standard deviation is 0.07. The correlation length for the initial saturation is still $50km$, and the initial saturation profiles are again hydrostatic.

The structure of the model error is the same as in Reference Experiment I, but the standard deviations are drastically increased. The standard deviation for the top moisture flux condition is now the equivalent of $50W/m^2$ in latent heat flux. Likewise, the standard

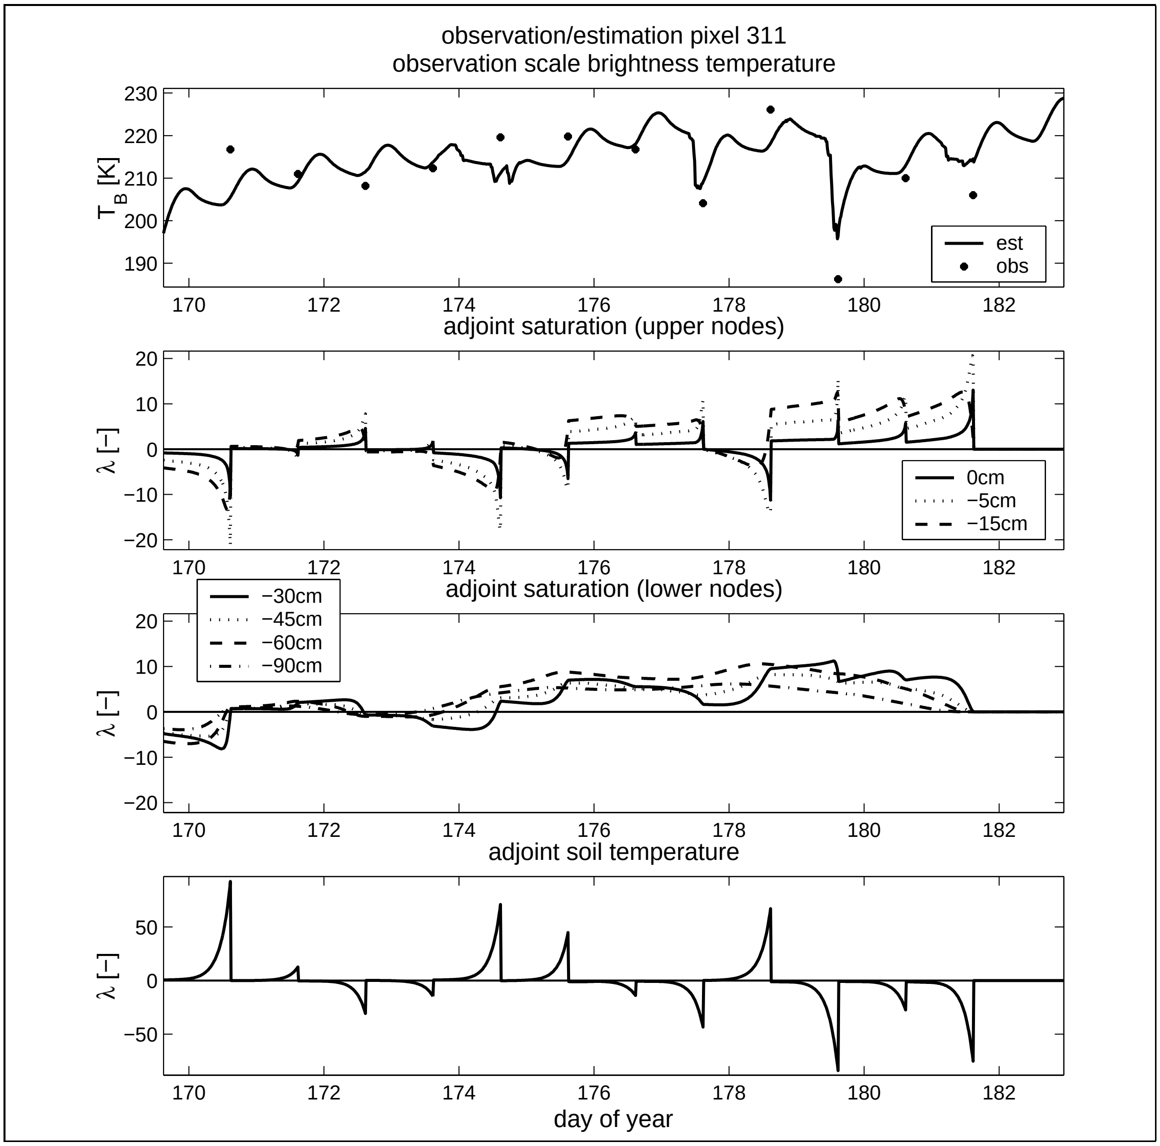

Figure 6.10: Adjoint variables for Reference Experiment I. The top panel shows the estimated brightness temperature for a representative observation pixel along with the (synthetic) observations. Note that for this experiment, observation and estimation pixels coincide. The second and the third panels show the adjoint saturation for the upper three nodes and the lower four nodes, respectively. The last panel shows the adjoint soil temperature. The adjoint equation is solved backward in time starting from zero at the final time (day 183). It is forced at observation times with the misfit between the observations and the estimates of the measured brightness temperatures. If the observed brightness temperature is higher than the estimated brightness, the saturation adjoint variable will be forced towards negative values, because the brightness temperature increases with decreasing saturation. The opposite is true for the adjoint soil temperature. Note how the upper two nodes of the adjoint saturation at $0cm$ and $-5cm$ are directly and instantaneously influenced by the data misfit forcing, while the rest of the nodes experience the forcing with a delay when the shock is propagated through the model physics.

deviation of the forcing terms in the force-restore equation and in the canopy energy balance equation is set to $50 W/m^2$. Moreover, the horizontal correlation of the model error has been increased to $6km$. This implies that the model errors are correlated over a few neighboring pixels. This correlation increases the share of the process noise update to approximately 25% of the total CPU time (Section 8.1.2).

From the new realizations of the uncertain parameters we again derive (synthetic) true fields and (synthetic) brightness data. The $5km$ resolution of the estimation and the observation pixels is unchanged, and we again use twelve brightness images (Figure 6.2 and Section 6.1).

6.2.1 Estimation of the True Fields

Figure 6.11 shows the true, the prior, and the estimated top node saturation for Reference Experiment II at six different times during the two-week period. As in Reference Experiment I, the soil moisture estimates are greatly improved through the assimilation of the brightness data. In Figure 6.12 we plot the area average errors in the top node saturation with respect to the true fields. The time and area average error is now 2.9% in saturation, compared to 1.4% in Reference Experiment I. This loss of quality in the estimate can be attributed to the much stronger model errors.

The strong model error is also responsible for the rapid increase of the error in the estimates at the end of the assimilation window. The last brightness image is assimilated on day 181.6 (Figure 6.2). After this time, we essentially forecast the top node saturation, which is of course very difficult to do with such strong model errors. Note that for this experiment the area average prior error does not decrease over the two-week period. Since we continuously add strong model error, the prior fields will not converge artificially to the truth. Reference Experiment II therefore describes a much more realistic scenario.

The lower panel of Figure 6.12 shows the area average root-mean-square errors of the prior and the estimated soil temperature. The stronger model error in the force-restore equation leads to a higher area average prior soil temperature error of $1.6K$ compared to $0.37K$ in Reference Experiment I. Moreover, the estimate is only a slight improvement over the prior guess, because L-band brightness temperature observations once a day are too infrequent compared to the short memory of soil temperature.

Figure 6.13 examines the estimate of the model errors. For a representative pixel, we plot the estimate, the prior, and the true model error time series. For comparison, we also show the 24 hour moving average of the true model error. Recall that the prior model error is zero. The top panel of Figure 6.13 shows the model error estimate for the moisture flux upper boundary condition. While the true model error correlation time is 10 hours, we only assimilate brightness updates once per day. Therefore the estimate is necessarily much smoother than the true time series. But if we compare the estimate to the 24 hour moving average of the true model error, we see that the estimate is in fact quite good.

As shown in the second panel of Figure 6.13, we get a poorer estimate for the model error in the soil energy balance. For the model error in the canopy energy balance, shown in the third panel, we do even worse. Obviously, the daily L-band brightness images do not contain much information on the canopy states. Although we have only plotted the results for a single pixel, we reach the same conclusion when we look at all pixels. Finally note that all model error estimates vanish after the last observation time plus one correlation time (day 181.6 plus 10 hours). The observations do not contain any information on times

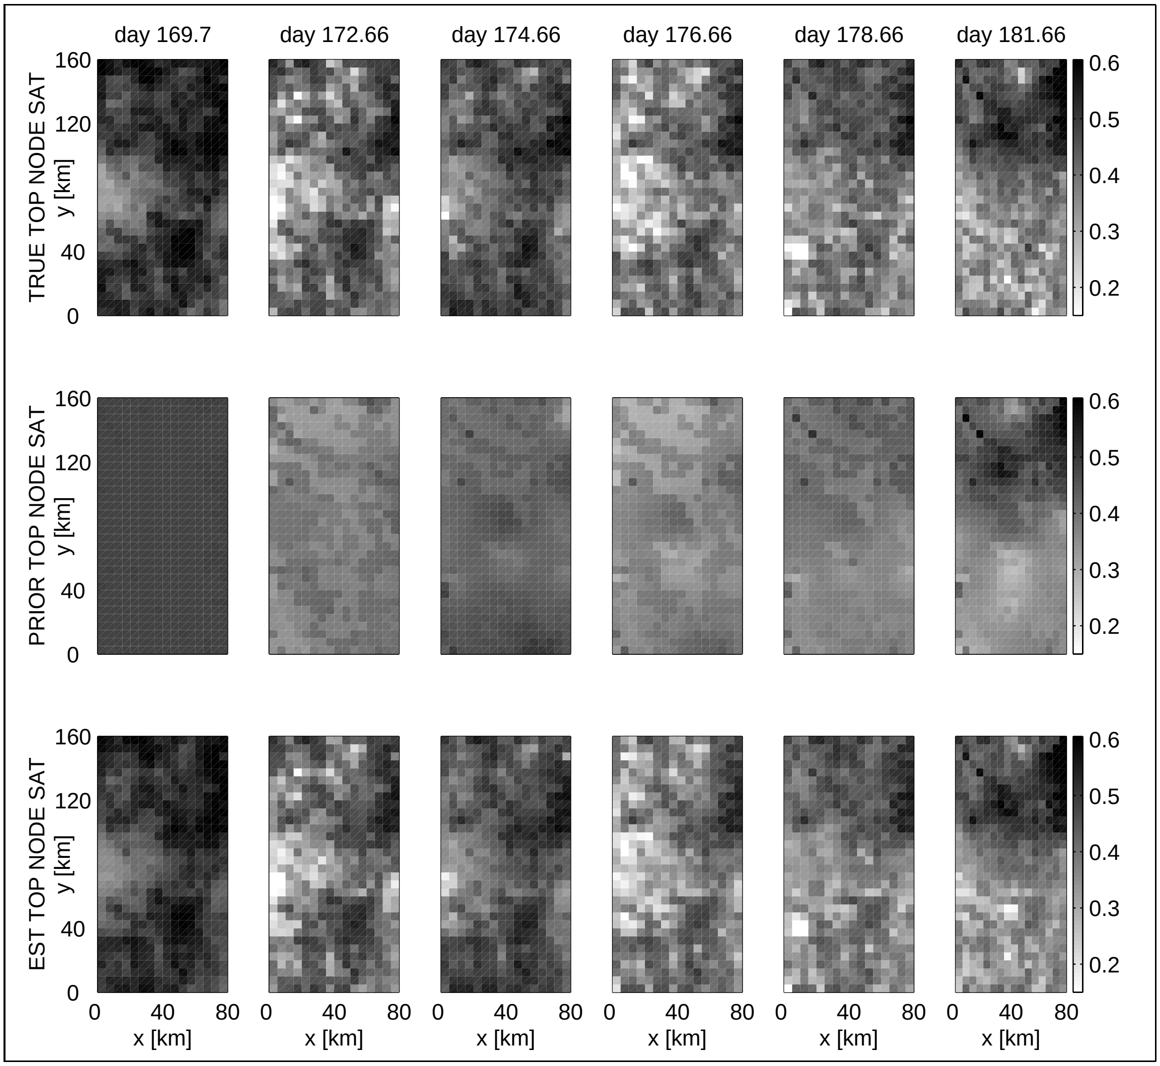

Figure 6.11: Top node saturation for Reference Experiment II. The first row shows the true top node saturation at six different times during the assimilation interval. In the second and third rows the prior and the estimate of the top node saturation, respectively, are depicted for the same six times.

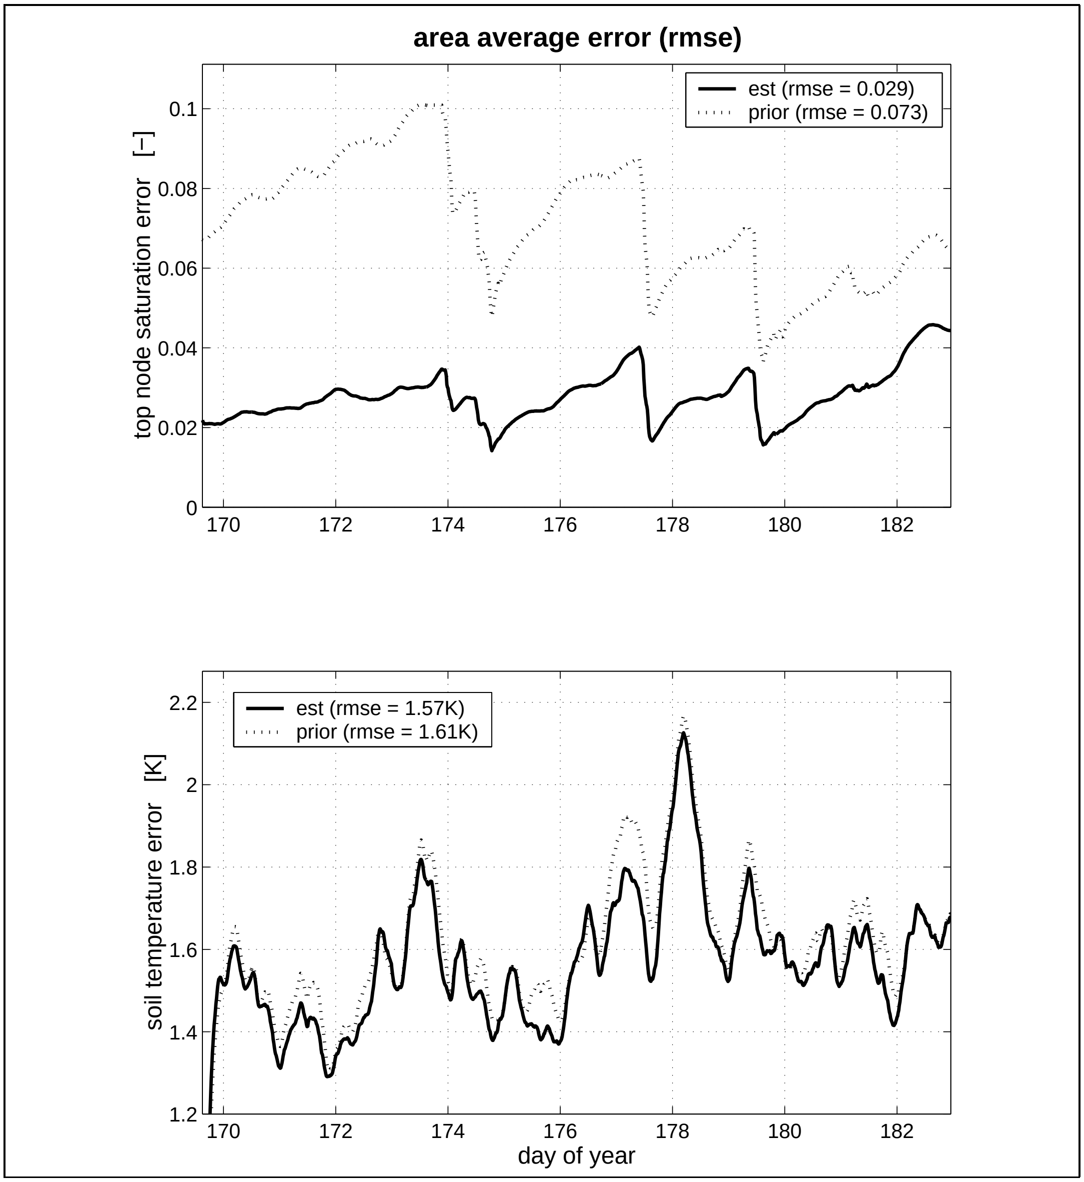

Figure 6.12: Area average errors for Reference Experiment II. The root-mean-square errors (rmse) of the prior and the estimated top node saturation and soil temperature with respect to the (synthetic) true fields are shown. In the legend we also indicate the temporal average of the area average rmse. Note that the soil moisture errors are in terms of saturation. Obviously, the assimilation greatly improves the errors in the top node saturation over the prior errors. On the other hand, daily brightness observations are not enough to improve the prior guess of the soil temperature by much.

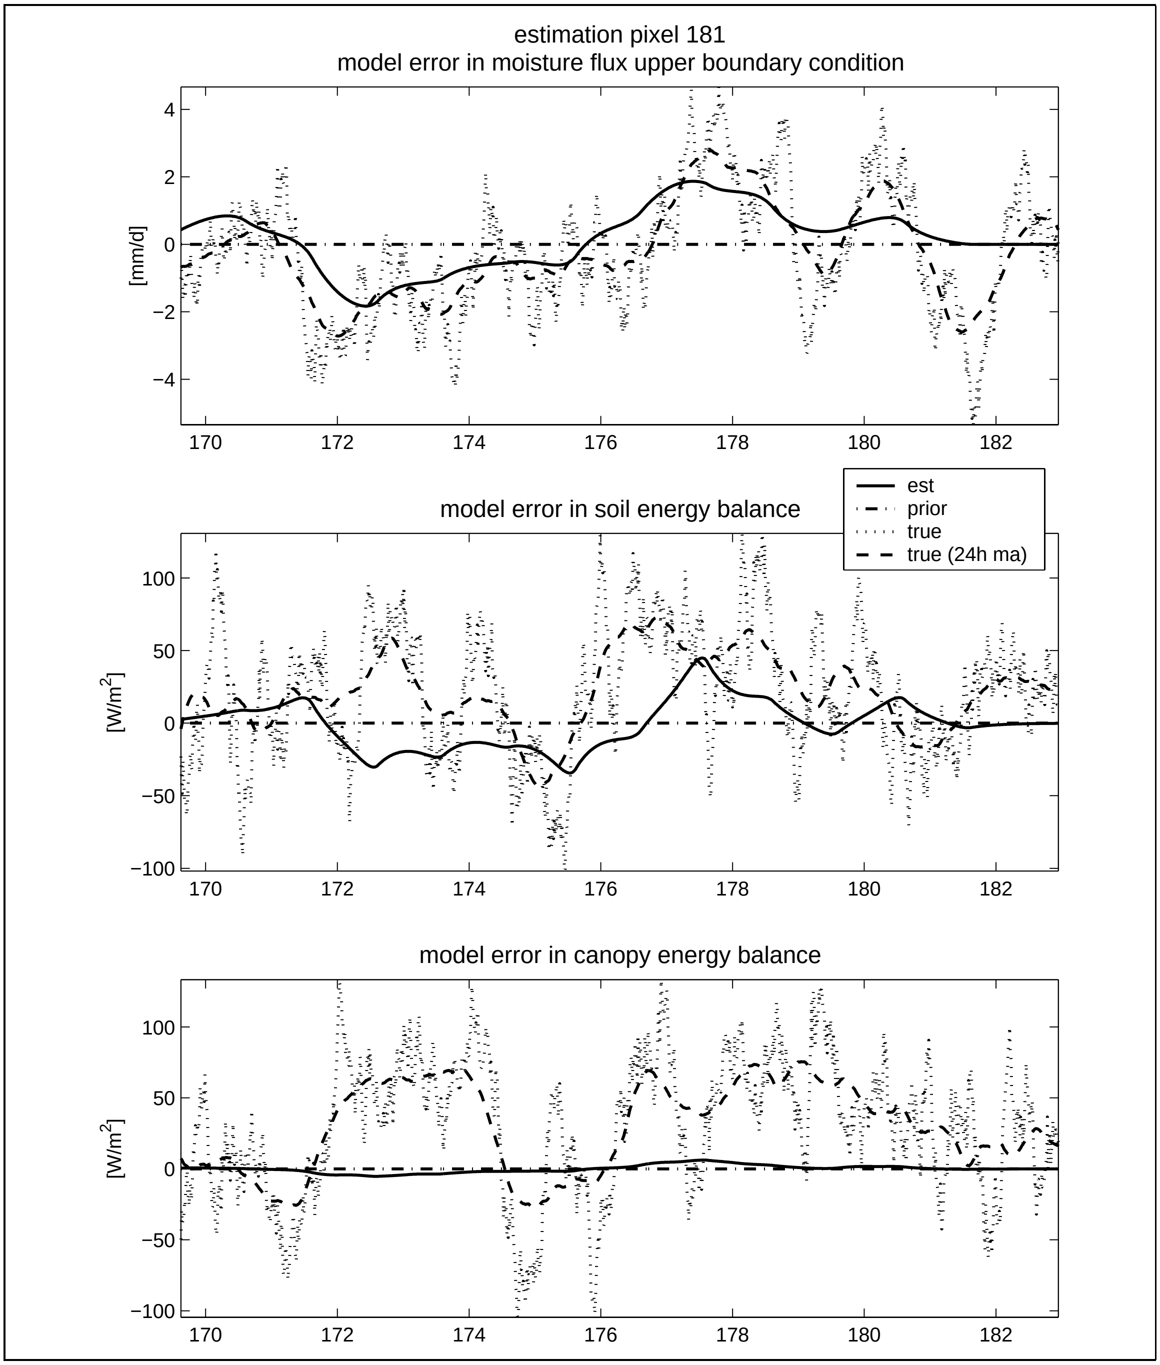

Figure 6.13: Model errors for Reference Experiment II. The estimated, the prior, and the true model errors for a representative pixel are shown. For comparison, we also plot the 24 hour moving average of the true time series. The upper panel shows the model error in the upper moisture flux boundary condition. Considering the difference in time scales between the true model error variability and the frequency of the brightness data, we get a very good estimate. In contrast, we have somewhat less skill in estimating the model error in the soil energy balance (second panel). We do even worse on the model error in the canopy energy balance (third panel).

that are further into the future.

For the subsurface nodes, we again get the same excellent estimates as for the surface node. To illustrate this point, Figure 6.14 shows the estimated, the prior, and the true profile saturation for a representative pixel. However, the limitations discussed in Section 6.1.3 still apply. In summary, we get an overly optimistic profile estimate because we use a fixed shape for the initial profile and because the assimilation window is too short for significant divergence to occur between the shape of the prior and the true profiles.

6.2.2 Reduced Objective Function and Posterior Data Residuals

Figure 6.15 shows the reduced objective function versus the iteration number. The converged value of the reduced objective is 5985 and lies within 1.5 standard deviations from the expected value of 6144. From the reduced objective function we have therefore no indication that the assimilation was not optimal.

Finally, a close look at the posterior data residuals yields qualitatively the same results as for Reference Experiment I (Figures 6.7, 6.8, and 6.9). The raw mean for all residuals with a 95% confidence interval is $0.04 \pm 0.11K$. For all but one of the residual brightness images we find a mean whose 95% confidence interval includes zero. The slight bias observed in the Reference Experiment I is absent here. The residuals are uncorrelated in space and in time, but the residuals of the individual images are not exactly normally distributed. We can again attribute this deviation to the nonlinear nature of the problem.

6.3 Downscaling Experiments

We now investigate the downscaling capability of the assimilation algorithm. As outlined in Section 4.7, we can effectively increase the resolution of the brightness images by making use of the fact that the inputs to the hydrologic model are available at a finer scale. The additional information is implicitly deduced from our knowledge of soil and land cover parameters as well as from the meteorologic forcings.

In the downscaling experiments, we use the same setup as in Reference Experiment I. In particular, we continue to estimate the land surface states at $5km$ resolution. In contrast to the reference setup, we now generate brightness observations at resolutions of $10km$ and $20km$, respectively. In the former scenario, each observation pixel contains four estimation pixels. Consequently, we call this setup the (1:4) downscaling scenario. In the case of brightness observations at $20km$ resolution, there are 16 estimation pixels within each observation pixel, and we call this setup the (1:16) downscaling scenario.

6.3.1 Estimation of the True Fields

In Figure 6.16, we compare the estimates of the top node saturation for both downscaling scenarios. For comparison, the true and the prior top node saturation are shown in the first two columns. All columns depict the saturation fields for the same three times during the experiment. For each downscaling estimate, we also show the outline of the corresponding observation pixels.

In both downscaling scenarios, the algorithm can adequately estimate the large-scale spatial distribution of the saturation. More importantly, structures at scales well below the scale of the observations can be resolved satisfactorily (Figure 6.16). This means that

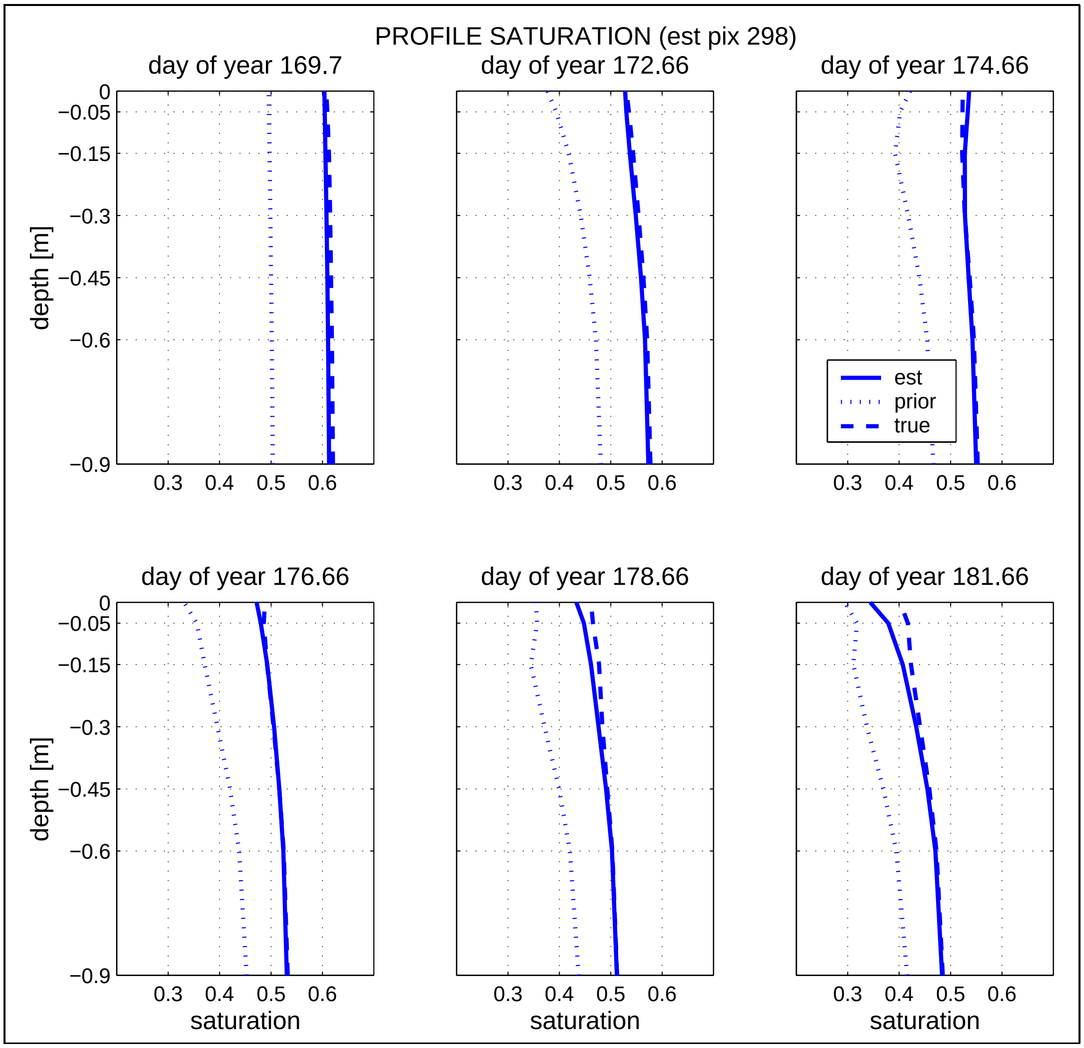

Figure 6.14: Profile saturation for Reference Experiment II. The true, the prior, and the estimated profile saturation are shown at six different times during the assimilation interval. Note that the profile estimates are overly optimistic because of the setup of the synthetic experiment.

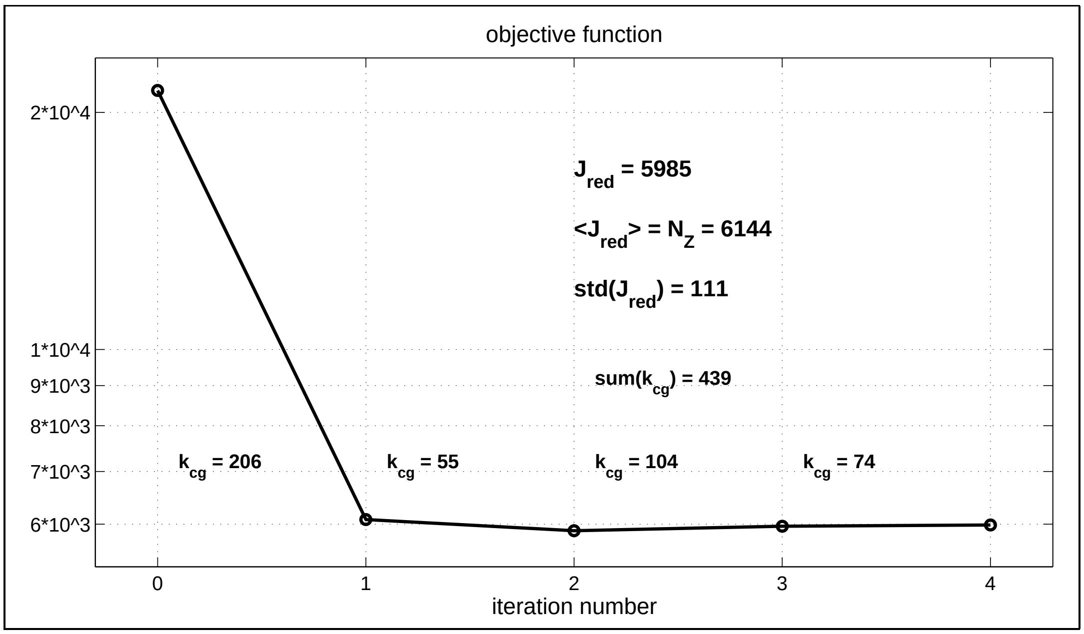

Figure 6.15: Objective function versus iteration number for Reference Experiment II. The reduced objective function after convergence is 5985. The number of data points is 6144, which is also the expected value of the reduced objective function. The standard deviation of the reduced objective function is 111. The values of k_{cg} indicate the number of linear combinations of representer functions that needed to be evaluated during the conjugate gradient iteration of the indirect representer approach (Chapter 8).